A complete listing of the pending claims is provided below.

1.   (Previously Presented) A method of communicating data between a plurality of processing nodes within an array of processor nodes, the method comprising:

determining a route for a unidirectional channel from a source processing node to a destination processing node within the array of processor nodes, the determined route based on a physical description of the array of processor nodes;

generating the unidirectional channel along the determined route from the source processing node to the destination processing node, the channel having a bandwidth requirement;

accepting the channel in the destination processing node;

allocating a transmit buffer for the channel in the source processing node;

allocating a receive buffer for the channel in the destination processing node;

writing data from a source processing element to the transmit buffer;

transmitting the data from the transmit buffer over the channel using a source network interface in the source processing node;

receiving the data over the channel into the receive buffer using a destination network interface in the destination processing node; and

reading the data from the receive buffer into the destination processing element.

2.   (Previously Presented) The method of claim 1 wherein the unidirectional channel is associated with a first task executing on the source processing element and a second task executing on the destination processing element.

3.   (Previously Presented) The method of claim 1 wherein the unidirectional channel is associated with a first port in the source processing element and a second port in the destination processing element.

4.   (Previously Presented) The method of claim 1 wherein the unidirectional channel has a maximum number of buffers and size of buffers.

5.   (Previously Presented) The method of claim 1 further comprising reserving intermediate resources for the unidirectional channel based on the bandwidth requirements.

6.   (Original) The method of claim 1 further comprising guaranteeing bandwidth based on the bandwidth requirements using time division multiplexing.

7.   (Original) The method of claim 1 further comprising guaranteeing bandwidth based on the bandwidth requirements using spatial division multiplexing.

8.   (Previously Presented) The method of claim 1 further comprising polling a plurality of channels to check if data is received into the receive buffer for the unidirectional channel.

9.   (Previously Presented) The method of claim 1 further comprising freeing the transmit buffer using the source processing element.

10.   (Previously Presented) The method of claim 1 further comprising freeing the receive buffer using the destination processing element.

11.   (Previously Presented) The method of claim 1 further comprising destroying the unidirectional channel.

12.   (Previously Presented) The method of claim 1 further comprising receiving a pointer for the data in the receive buffer into the destination processing element and wherein reading the data from the receive buffer into the destination processing element is based on the pointer.

13.   (Original) The method of claim 1 wherein a time for a receive call in the destination processing element does not depend upon a size of the data.

14.   (Previously Presented) A multi-processor system comprising:

a source processing node, comprising,

   a source processing element configured to generate a unidirectional channel,
      allocate a transmit buffer for the unidirectional channel, and write
      data to the transmit buffer for the unidirectional channel, and

   a source network interface configured to transmit the data from the transmit
      buffer of the source processing node over the unidirectional channel;
      and

a destination processing node comprising,

   a destination processing element configured to accept the unidirectional
      channel, allocate a receive buffer for the unidirectional channel in the
      destination processing node, and receive the data from the receive
      buffer,

   a destination network interface configured to receive the data into the receive
      buffer for the unidirectional channel,

   the unidirectional channel having a bandwidth requirement and generated
      by the source processing element along a route, the route and the
      bandwidth requirement based on one or more tasks associated with
      the destination processing node, and

   the source processing node and destination processing node included within
      an array of processing nodes.

15.   (Previously Presented) The multi-processor system of claim 14 wherein the
unidirectional channel is associated with a first task executing on the source processing
element and a second task executing on the destination processing element.

16.   (Previously Presented) The multi-processor system of claim 14 wherein the
unidirectional channel is associated with a first port in the source processing element and a
second port in the destination processing element.

17.   (Previously Presented) The multi-processor system of claim 14 wherein the unidirectional channel has a maximum number of buffers and size of buffers.

18.   (Previously Presented) The multi-processor system of claim 14 wherein the source processing node and the destination processing node are configured to reserve intermediate resources for the unidirectional channel based on the bandwidth requirements.

19.   (Original) The multi-processor system of claim 14 wherein the source processing node is configured to guarantee bandwidth based on the bandwidth requirements using time division multiplexing.

20.   (Original) The multi-processor system of claim 14 wherein the source processing node is configured to guarantee bandwidth based on the bandwidth requirements using spatial division multiplexing.

21.   (Previously Presented) The multi-processor system of claim 14 wherein the destination processing element is configured to poll a plurality of channels to check if data is received into the receive buffer for the unidirectional channel.

22.   (Original) The multi-processor system of claim 14 wherein the source processing element is configured to free the transmit buffer.

23.   (Original) The multi-processor system of claim 14 wherein the destination processing element is configured to free the receive buffer.

24.   (Previously Presented) The multi-processor system of claim 14 wherein the source processing element is configured to destroy the unidirectional channel.

25.   (Original) The multi-processor system of claim 14 wherein the destination processing element is configured to receive a pointer for the data in the receive buffer into the destination processing element and receive the data from the receive buffer based on the pointer.

26.   (Original) The multi-processor system of claim 14 wherein a time for a receive call in the destination processing element does not depend upon a size of the data.

27.   (Withdrawn) A method of compiling applications for a multi-processor system, the method comprising:

receiving a physical description of the multi-processor system;

receiving an application description indicating tasks for the applications and channels for communications between the tasks;

processing the physical description and the application description to determine routing information for the channels and to assign the tasks to processors in the multi-processor system; and

generating executable code for the processors based on the physical description and the application description.

28.   (Withdrawn) The method of claim 27 wherein the physical description and the application description are in a package description.

29.   (Withdrawn) The method of claim 27 wherein the physical description includes a configuration of processors in the multi-processor system.

30.   (Withdrawn) The method of claim 27 wherein the application description includes application code for the tasks.

31.   (Withdrawn) The method of claim 27 wherein the application description includes assignments of the tasks to execute on processors of the multi-processor system.

32.   (Withdrawn) The method of claim 27 wherein the application description includes channels for communications.

33.   (Withdrawn) The method of claim 32 wherein the application description includes routing for the channels.

34.   (Withdrawn) The method of claim 27 wherein the application description includes shared memory descriptions.

35.   (Withdrawn) The method of claim 27 further comprising processing the physical description and the application description to check for syntax and semantic errors.

36.   (Withdrawn) The method of claim 27 wherein processing the physical description and the application description to determine routing information for the channels further comprises generating routing tables for the channels.

37.   (Withdrawn) The method of claim 27 further comprising generating boot code for the processors in the multi-processor system.

38.   (Withdrawn) The method of claim 27 further comprising setting scheduling policies for tasks executing in the processors.

39.   (Withdrawn) The method of claim 27 wherein the executable code is for each processor based on a processor number.

40.   (Withdrawn) The method of claim 27 further comprising mapping the executable code to memory in the multi-processor system.

41.   (Withdrawn) The method of claim 27 wherein generating the executable code further comprises linking boot code, operating system code, and application code for the tasks.

42.   (Withdrawn) The method of claim 27 wherein the executable code is executed on a host system for emulation.

43.   (Withdrawn) The method of claim 27 wherein the executable code is executed on a simulator.

44.   (Withdrawn) The method of claim 27 further comprising generating source code of the executable code for debugging.

45.   (Withdrawn) The method of claim 27 further comprising:
        determining assignment of the tasks in the multi-processor system based on the physical description;
        determining the channels for communications between the tasks; and
        generating the application description based on task assignments and the channels.

46.   (Withdrawn) The method of claim 27 further comprising:
        receiving performance data based on the execution of the executable code;
        generating the application description based on the performance data.

47.   (Withdrawn) A software product for compiling applications for a multi-processor system, the software product comprising:
        package compiler software operational when executed by a first processor to direct the first processor to receive a physical description of the multi-processor system, receive an application description indicating tasks for the applications and channels for communications between the tasks, process the physical description and the application description to determine routing information for the channels and to assign the tasks to

second processors in the multi-processor system, and generate executable code for the second processors based on the physical description and the application description; and

a software storage medium operational to store the package compiler software

48. (Withdrawn) The software product of claim 47 wherein the physical description and the application description are in a package description.

49. (Withdrawn) The software product of claim 47 wherein the physical description includes a configuration of the second processors in the multi-processor system.

50. (Withdrawn) The software product of claim 47 wherein the application description includes application code for the tasks.

51. (Withdrawn) The software product of claim 47 wherein the application description includes assignments of the tasks to execute on the second processors of the multi-processor system.

52. (Withdrawn) The software product of claim 47 wherein the application description includes channels for communications.

53. (Withdrawn) The software product of claim 52 wherein the application description includes routing for the channels.

54. (Withdrawn) The software product of claim 47 wherein the application description includes shared memory descriptions.

55. (Withdrawn) The software product of claim 47 wherein the package compiler software is operational when executed by the first processor to direct the first processor to

process the physical description and the application description to check for syntax and semantic errors.

56.   (Withdrawn) The software product of claim 47 wherein the package compiler software is operational when executed by the first processor to direct the first processor to generate routing tables for the channels.

57.   (Withdrawn) The software product of claim 47 wherein the package compiler software is operational when executed by the first processor to direct the first processor to generate boot code for the second processors in the multi-processor system.

58.   (Withdrawn) The software product of claim 47 wherein the package compiler software is operational when executed by the first processor to direct the first processor to set scheduling policies for tasks executing in the second processors.

59.   (Withdrawn) The software product of claim 47 wherein the executable code is for each second processor based on a processor number.

60.   (Withdrawn) The software product of claim 47 wherein the package compiler software is operational when executed by the first processor to direct the first processor to map the executable code to memory in the multi-processor system.

61.   (Withdrawn) The software product of claim 47 wherein the package compiler software is operational when executed by the first processor to direct the first processor to link boot code, operating system code, and application code for the tasks to generate the executable code.

62.   (Withdrawn) The software product of claim 47 wherein the executable code is executed on a host system for emulation or simulation.

63.   (Withdrawn) The software product of claim 47 wherein the executable code is executed on a simulator.

64.   (Withdrawn) The software product of claim 47 wherein the package compiler software is operational when executed by the first processor to direct the first processor to generate source code of the executable code for debugging.

65.   (Withdrawn) The software product of claim 47 wherein the package compiler software is operational when executed by the first processor to direct the first processor to determine assignment of the tasks in the multi-processor system based on the physical description, determine the channels for communications between the tasks, and generate the application description based on task assignments and the channels.

66.   (Withdrawn) The software product of claim 47 wherein the package compiler software is operational when executed by the first processor to direct the first processor to receive performance data based on the execution of the executable code and generate the application description based on the performance data.

67.   (Withdrawn) A method of booting a multi-processor system comprising a root processor and at least one non-root processor, the method comprising:

identifying the root processor in the multi-processor system that does not have memory associated with the at least one non-root processor;

transmitting a boot message from the root processor to the at least one non-root processors;

receiving the boot message into the at least one non-root processor;

in the at least one non-root processor, obtaining the non-root boot code based on the message in the non-root code; and

configuring the at least one non-root processor based on the non-root boot code.

68.   (Withdrawn) The method of claim 67 further comprising executing processor initialization code for initialization of the root processor.

69.   (Withdrawn) The method of claim 67 further comprising executing processor initialization code for initialization of the at least one non-root processor.

70.   (Withdrawn) The method of claim 67 wherein the boot message is a Joint Test Action Group command.

71.   (Withdrawn) The method of claim 67 wherein transmitting the boot message is from a Joint Test Action Group port.

72.   (Withdrawn) The method of claim 67 wherein receiving the boot message is into a Joint Test Action Group port.

73.   (Withdrawn) The method of claim 67 further comprising:

in the at least one non-root processor, transmitting a boot complete message to the root processor;

in the root processor, transmitting a proceed message in response to the boot complete message; and

in the at least one non-root processor, beginning operations in response to the proceed message.

74.   (Withdrawn) A multi-processor system comprising:

a root processor configured to identify the root processor as a root and transmit a boot message from the root processor to the at least one non-root processors;

the at least one non-root processor that does not have memory associated with the at least one non-root processor and that is configured to receive the boot message, obtain the non-root boot code based on the message in the non-root code, and configuring the at least one non-root processor based on the non-root boot code.

75.   (Withdrawn) The multi-processor system of claim 74 wherein the root processor is configured to execute processor initialization code for initialization of the root processor.

76.   (Withdrawn) The multi-processor system of claim 74 wherein the at least one non-root processor is configured to execute processor initialization code for initialization of the at least one non-root processor.

77.   (Withdrawn) The multi-processor system of claim 74 wherein the boot message is a Joint Test Action Group command.

78.   (Withdrawn) The multi-processor system of claim 74 wherein the root processor is configured to transmit the boot message from a Joint Test Action Group port.

79.   (Withdrawn) The multi-processor system of claim 74 wherein the at least one non-root processor is configured to receive the boot message into a Joint Test Action Group port.

80.   (Withdrawn) The multi-processor system of claim 74 wherein the at least one non-root processor is configured to transmit a boot complete message to the root processor and begin operations in response to a proceed message and wherein the root processor is configured to transmit the proceed message in response to the boot complete message.

81.   (Previously Presented) The method of claim 1, further comprising:

receiving the first task in the source processing node, wherein the step of generating a channel is performed in response to receiving the first task.

82   (Previously Presented) The method of claim 1, further comprising:

determining a topology of processing nodes to process one or more tasks, the topology including the channel.

83   (Previously Presented) The method of claim 1, wherein the step of accepting the unidirectional channel includes:

receiving a response signal from the destination processing node by the source processing node.

84.   (Previously Presented) The method of claim 1, further comprising:

assigning tasks to one or more nodes in the array of nodes, wherein said step of generating the unidirectional channel is performed in response to said step of assigning tasks.

85.   (Previously Presented) The multi-processor system of claim 14, wherein the route for the unidirectional channel is further based on a physical description of the multi-processor system.

86.   (Previously Presented) The multi-processor system of claim 14, wherein a compiler is configured to determine routing information for one or more channels and assign a task to one or more destination processing nodes.

87.   (Previously Presented) The multi-processor system of claim 14, wherein the transmit buffer and receive buffer are allocated based on the one or more tasks and a physical description of a portion of the array of nodes over which the data is to be transmitted.

88   (Previously Presented) The method of claim 1, wherein determining a route for a unidirectional channel is based on an application description.

89   (Previously Presented) The method of claim 1, wherein generating the unidirectional channel along the determined route is based on an allocated communication bandwidth between a first task on the source processing node and a second task on the destination processing node.